A Bridge of Bits

Norman Margolus
Laboratory for Computer Science
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
Cambridge, MA 02139

Abstract

Cellular Automata are discrete physics-like systems that can be exactly simulated by digital hardware. These systems can incorporate many realistic physical constraints and still be capable of performing digital computation. Such systems bridge the boundary between physical models and computational models, and so can play a central role in investigating the relevance of physical ideas to the theory and practice of computation, and computational ideas to the construction and use of physical models.

1 Introduction

Physics is the ultimate machine language. It is the constraints and opportunities provided by the laws of nature that ultimately determine what can be computed, and how, by our engines that transform information. As ever increasing demands for greater computational power and efficiency drive us into increasingly microscopic realms technologically, we are forced to adapt both our models of computation and our algorithms to fit the physics.

Anticipating these trends leads us naturally to the study of Cellular Automata (CA): discrete physics-like systems that can be exactly simulated by digital hardware. By incorporating physical constraints such as locality of interaction, three-dimensionality of interconnection, and uniformity of structure, such systems act as laboratories within which we can study the possibilites of digital computation in the face of microphysical constraints.

Because of the massive fine-grained parallelism that is possible for physical realizations of CA, theoretical progress in making efficient use of CA models has direct practical consequences. But perhaps even more interesting are the theoretical issues that arise when we strive to bring our fundamental models of computation and of physics as close together as possible.

2 Why bring the models together?

CA are representative of a class of models in which we try to incorporate computational features in a context that is compatible with physical structures and constraints. There are several reasons to investigate such informational models which emulate physics.

- *Efficiency:* I have sometimes likened the process of refining our computers to that of refining an ore. We could imagine that we are trying to ultimately create a chunk of 100% pure *Computronium*, in which as little as is physically possible is waste-material. Every available degree of freedom is being used in a useful computational mode, with both interaction sites and communication paths being as small and dense and fast as possible. To achieve such an end, our computation would have to map almost exactly onto the structure and dynamics of the matter: the computation would look like the microscopic dynamics of an extended material system. Thus studying models of computation which incorporate realistic physical structure and constraints is a step towards harnessing the massive fine-grained parallelism inherent in the laws of nature.

- *A Science of Computation:* Computer science doesn't presently have the kinds of strong principles and global techniques that physics does. In the coming age of efficient computers with a physics-like structure, particularly in physical simulations, much of the conceptual and analytic machinery of physics will be applicable to computers.

- *Physics modeling:* Traditional physical models were not conceived with digital simulation in mind. More recently, original computational models of physical systems (such as lattice gases[10]) have been developed which map much more directly onto digital hardware (see Fig. 1).

0-8186-3420-0/93 $03.00 © 1993 IEEE

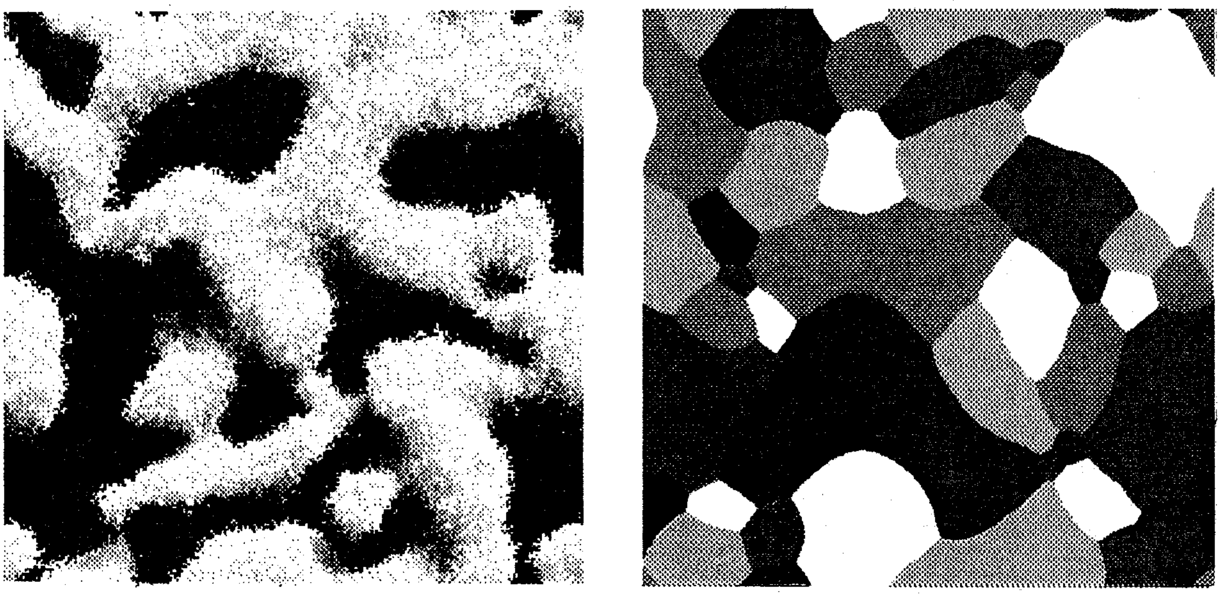

Figure 1: Two physical simulations done on CAM-8. (a) Spongy three-dimensional structure obtained by "majority" annealing. (b) Typical attractor texture produced by one of David Griffeath's "plurality" rules.

> Adapting our models to the computer will become an increasingly important activity, since we are free to change our models, but not the ultimate hardware constituted by the laws of physics. I think that we will find that informational models of physics are not only easier to simulate, but also ultimately more fundamental.

Thus there are practical as well as theoretical motivations for finding and investigating models which bring physics and computation closer together. Not only may they lead to maximally efficient computation, but they can act as bridges between ideas in physics and computation, allowing both concepts and techniques to be transferred between the two disciplines.

3 Cellular automata

Cellular automata are models that have one foot in the realm of computation, and the other foot in the realm of physics. The following list summarizes some of the properties from each of these two realms that have been incorporated simultaneously into CA models; we also list some properties that are noticeably missing from our CA models.

- **Computer realm:**
  Digital
  Exact
  Universal

- **Physics realm:**
  Space
  Time
  Locality of interaction
  Finite speed of information propagation
  3-dimensional interconnectivity
  Uniform laws
  Conservation laws
  Microscopic reversibility
  Finite entropy for a finite system
  Entropy proportional to volume

- **Missing:**
  Relativity
  QM

From the computer world we take the digital abstraction, which allows a real physical system to operate for an indefinite number of steps without accumulating any error. This permits hardware to simulate a model exactly. We also add the property of universality: a universal CA can be configured to exactly simulate the operation of any physically possible information process.[1] This can be done by changing only the initial state of our CA without changing the underlying updating law, much as we build machines in the physical world without changing the underlying laws of physics.

From nature, we abstract a number of properties that we put into our CA. First we include discrete versions of space and time, with 3-dimensional interconnectivity and a finite-range (neighborhood) interaction. As in nature, the laws of our discrete world are the same everywhere. Since information can move at most the interaction-range distance in one time-step of the system, this defines a maximum speed of information propagation.

Other properties are included mainly by putting constraints on the law and the neighborhood it uses. For example, there are several simple techniques available to achieve conservation laws and reversibility (conservation of information) in CA systems. Perhaps the simplest involves partitioning: we divide all of the data bits[2] in our space into disjoint local groupings, and then update each group independently of all others, replacing all of the bits in a given group with new values. We then rearrange the bits between neighboring groups, and repeat the process. Now observe: if our law for the replacement is invertible, then the overall dynamics will be invertible. If our law conserves the number of ones in each group, then the number of ones in our space will be a conserved quantity.

Reversible CA (RCA) are particularly physics-like, because they support a realistic thermodynamics[6, 14]; computational structures that "live" inside RCA are constrained to output (on the average) at least as much "entropy" as comes into them, since they can't invertibly erase information. Thus entropy becomes an important concept in RCA, and the connection between microscopic reversibility and the second law of thermodynamics is illustrated in a simple system.

Entropy in RCA models is an extensive quantity. If, for example, we use the presence of a 1 in a cell of an RCA to represent the presence of a particle (in a lattice gas simulation, say), then interchanging two such "identical particles" (1's) doesn't result in a new configuration. Thus, in general, the Gibbs paradox is automatically avoided in RCA models.

This example also illustrates the fact that there is nothing particularly quantum mechanical about the concept of identical particles—they are basic to any informational model. The finiteness of the entropy of a finite system is also automatically built into such discrete models; this is also a property that, in physics, derives from quantum mechanics.

Finally, we have seen that some aspects of both relativity (finite speed of information propagation) and quantum mechanics (identical particles, finite-state for a finite system) are built into our CA models. But there are important aspects of these fundamental physical models which have not (as yet) been reconciled with computational models.

4 Relativity in CA

There are a number of interesting outstanding questions regarding how much like real physics CA worlds can be. One of the most basic, to my mind, concerns whether such fully discrete and finite models can accommodate (in an appropriate macroscopic limit) fully relativistic behavior of composite systems, along with exact microscopic reversibility and universal computing capability. For contrast, consider for a moment the well known (irreversible) Game Of Life. This is a universal CA, and so is capable of supporting arbitrary complexity. But there is no real notion in a Life "universe" of the "same" composite system in various states of motion. In a (macroscopically) relativistically invariant CA world, the same (macroscopic) laws of "physics" would apply to all inertially moving objects. The same complex organisms could exist in all states of motion, making such a world much more like ours than Life is (if composite structures can't move around, its hard to see how simple pieces can come together to form more complex objects). Of course, since Life is universal, once we found such a CA, we could simulate it with Life, but such a simulation would involve starting the whole Life universe off in some very special initial state.

Note that it is already well known how to achieve reversibility, universality, or relativistic invariance in CA models separately. My CA version[6] of Fredkin's Billiard Ball Model incorporates the first two constraints, while various finite difference schemes achieve relativistic invariance in the macroscopic limit. The real problem that we're addressing here is that of making computer models which simultaneously incorporate as many as possible of the fundamental constraints that real microscopic physical systems must obey.

[1] This universal simulation capability defines a universal computer. According to our best current understanding of the laws of physics, it is true that an ordinary digital computer, with enough memory, is universal in this sense. This is the kind of computer that our CA can simulate.

[2] Or trits, or whatever the data elements are.

4.1 Quantum Computation

So far our approach to the problem of bringing informational models and physics closer together has started mostly from the computer-model end: we have taken a digital dynamics and added physical properties and constraints to it. We could continue this approach, trying to simulate quantum mechanics with a CA model (the putatively impossible "hidden variables" problem[2]). Reference [24] describes some significant progress in this direction.

Here we will start instead from the physics end, beginning by asking if there are "reasonable" models of physical systems in which every physical degree of freedom is mapped onto a computational degree of freedom. Of course, it is precisely those computer models which incorporate fundamental physical constraints which are candidates for modeling with such "reasonable" physical models. In classical mechanics, Fredkin's Billiard Ball Model is a good example of this approach: he modeled reversible, bit-conserving logic using elastic collisions of billiard balls[5].

But our best models of nature are quantum mechanical, and so it makes sense to try to find microscopic QM models of computation[4, 9, 8, 12, 18]. Since the Schrödinger evolution of a closed QM system is unitary (and hence invertible), and since even in QM relativity requires microscopic physical dynamics to be *local* (locations that are spatially near each other affect each other most quickly), it is natural to think of trying to model a reversible CA as a QM system of coupled spins.

Before trying to make a more realistic model, one can begin by asking whether it is possible at all to have a model of the sort we have in mind, even given complete freedom to invent a suitable (hermitian and local) hamiltonian.[3] Choosing my CA version of the Billiard Ball Model as the computational system to model, I was able to achieve[18] (with some help from Mike Biafore) one dimension of parallelism in a two dimensional system of spins. This example pointed up that there are difficult (some now claim insurmountable) synchronization problems: it is hard to coordinate the actions of the different pieces that make up a coherent microscopic "QM" simulation of a fully parallel deterministic system in more than one dimension.

It turns out that the sorts of synchronization schemes one is driven to use to achieve deterministic parallel QM computation in one dimension involve only local synchrony: the system can be at widely different stages of the computation at widely separated positions. This suggests that a relativistically invariant CA would be more naturally modeled by such a system—this may be at least part of the problem. Or perhaps QM, at least as presently formulated, only admits deterministic fully-parallel computation at the macroscopic scale. If you equate such deterministic computation with classical mechanics (think of Fredkin's Billiard Ball Model) this seems like a real possibility. Thus by trying to bring QM and CA models close together, we are addressing questions about the essence of the difference between classical and quantum systems.

[3] Such *ideal* models are also directly useful for setting upper bounds on the physical resources a computation requires (e.g., energy tied up), as well as for studying fundamental issues of information in physics.

5 Cellular Automata Machines

To bring our CA models to life as miniature universes that can be watched and analyzed, Tom Toffoli and I have invested much of our time for several years on the design and construction of Cellular Automata Machines (CAMs)[13, 11]. More recently, I have completed the design and construction of a new kind of CAM[16, 20, 25], a Space Time Event Processor (STEP machine), that is particularly well suited to the simulation and analysis of CA systems that embody physical properties.

CAM-8 takes advantage of (1) the parallelism and uniformity of a CA space, and (2) a restriction to lattice-gas-like data movement (but with arbitrary velocities in n dimensions), to allow the hardware equivalent of a workstation to run CA models on many millions of sites several times faster than a Cray Y-MP or Thinking Machines CM-2. This is a modularly scalable architecture, with each module handling a chunk of space, and modules stacked together in 3-dimensions. The machine consists essentially of DRAMs used for cell states, and SRAMs used for lookup tables that implement the CA transition rule, with a little bit of custom glue logic to handle the data movement and tie everything together. There are no conventional processors.

We expect this machine to be an important adjunct to our more theoretical investigations of CA systems. We are also making arrangements to make copies of this machine available to the research community.

References

[1] LANDAUER, Rolf, "Irreversibility and Heat Generation in the Computing Process," *IBM J.* 5

(1961), 183–191.

[2] BELL, J. S., "On the Einstein Podolsky Rosen paradox," *Physics*, **1** (1964), 195–200; reprinted in *Quantum Theory and Measurement* (J. A. WHEELER and W. H. ZUREK eds.), Princeton University Press (1983).

[3] BENNETT, Charles, "Logical reversibility of computation," *IBM J. Res. Devel.* **6** (1973), 525–532.

[4] BENIOFF, P. A., "Quantum mechanical Hamiltonian models of discrete processes that erase their own histories: application to Turing machines," *Int. J. Theor. Physics* **21** (1982), 177–202.

[5] FREDKIN, Edward, and Tommaso TOFFOLI, "Conservative Logic," *Int. J. Theor. Phys.* **21** (1982), 219–253.

[6] MARGOLUS, Norman "Physics-like Models of Computation," *Physica* **10D** (1984), 81–95.

[7] TOFFOLI, Tommaso, "Cellular Automata as an Alternative to (Rather Than an Approximation of) Differential Equations in Modeling Physics," *Physica* **10D** (1984), 117–127.

[8] D. Deutsch, "Quantum theory, the Church-Turing hypothesis, and universal quantum computers," *Proc. Roy. Soc. Lond. A* **400** (1985), 97–117.

[9] R. P. Feynman, "Quantum mechanical computers," *Opt. News* **11** (1985).

[10] FRISCH, Uriel, Brosl HASSLACHER, and Yves POMEAU, "Lattice-Gas Automata for the Navier-Stokes Equation," *Phys. Rev. Lett.* **56** (1986), 1505–1508.

[11] MARGOLUS, Norman, Tommaso TOFFOLI, and Gérard VICHNIAC, "Cellular-Automata Supercomputers for Fluid Dynamics Modeling," *Phys. Rev. Lett.* **56** (1986), 1694–1696.

[12] MARGOLUS, Norman, "Quantum Computation," *New Techniques and Ideas in Quantum Measurement Theory* (Daniel GREENBERGER ed.), New York Academy of Sciences (1986), 487–497.

[13] TOFFOLI, Tommaso, and Norman MARGOLUS, *Cellular Automata Machines: A New Environment for Modeling*, MIT Press (1987).

[14] MARGOLUS, Norman, *Physics and Computation*, MIT Physics Doctoral Thesis (1987).

[15] BENNETT, Charles, Norman MARGOLUS and Tommaso TOFFOLI, "Bond-energy variables for spin-glass dynamics," *Phys. Rev. B* **37** (1988), 2254.

[16] MARGOLUS, Norman, and Tommaso TOFFOLI, "Cellular Automata Machines," *Lattice Gas Methods of Partial Differential Equations* (DOOLEN et al. eds.), Addison-Wesley Longman Publishing Group Ltd. (1989).

[17] TOFFOLI, Tommaso, "Frontiers in Computing," Information Processing (G. X. Ritter ed.), North-Holland (1989), 1.

[18] MARGOLUS, Norman, "Parallel Quantum Computation," *Complexity, Entropy, and the Physics of Information* (Wojciech ZUREK ed.), Addison-Wesley, 1990.

[19] TOFFOLI, Tommaso, and Norman MARGOLUS, "Invertible Cellular Automata: A Review," *Physica D*, 45(1-3), 1990.

[20] TOFFOLI, Tommaso, and Norman MARGOLUS, "Programmable matter," *Physica D*, 47:263-272, 1990.

[21] MARGOLUS, Norman, "Fundamental Physical Constraints on the Computational Process," to appear in the proceedings of the First Foresight Conference on Nanotechnology, MIT Press, 1991.

[22] MARGOLUS, Norman, *Physics-like Models of Computation*, MIT Press (1993). (This has not yet gone to press).

[23] SMITH, Mark, Yaneer BAR-YAM, Y. RABIN, N. MARGOLUS, T. TOFFOLI, and C. H. BENNETT, "Cellular Automaton Simulation of Polymers", *Complex Fluids* (D. WEITZ *et al.* ed.), Materials Research Society (1992).

[24] HRGOVČIĆ, H. J., "Quantum mechanics on spacetime lattices using path integrals in a Minkowski metric," MIT Physics Doctoral Thesis (1992); see also these proceedings.

[25] MARGOLUS, Norman, "Data-Array Computer and Processing Method," United States Patent and Trademark Office Patent No. 5,159,690, Filed 09/30/88, Issued 10/27/92.